

DOEF: A Dynamic Object Evaluation Framework

Zhen He¹ and Jérôme Darmont²

¹Department of Computer Science
The Australian National University
Canberra, ACT 0200, Australia
zhen@cs.anu.edu.au

²ERIC, Université Lumière Lyon 2
5 avenue Pierre Mendès-France
69676 Bron Cedex, France
jerome.darmont@univ-lyon2.fr

Abstract. In object-oriented or object-relational databases such as multimedia databases or most XML databases, access patterns are not static, i.e., applications do not always access the same objects in the same order repeatedly. However, this has been the way these databases and associated optimisation techniques like clustering have been evaluated up to now. This paper opens up research regarding this issue by proposing a dynamic object evaluation framework (DOEF) that accomplishes access pattern change by defining configurable styles of change. This preliminary prototype has been designed to be open and fully extensible. To illustrate the capabilities of DOEF, we used it to compare the performances of four state of the art dynamic clustering algorithms. The results show that DOEF is indeed effective at determining the adaptability of each dynamic clustering algorithm to changes in access pattern.

Keywords: Performance evaluation, Dynamic access patterns, Benchmarking, Object-oriented and object-relational databases, Clustering.

1 Introduction

Performance evaluation is critical for both designers of Object Database Management Systems (architectural or optimisation choices) and users (efficiency comparison, tuning). Note that we term Object Database Management Systems (ODBMSs) both object-oriented and object-relational systems, indifferently. ODBMSs include most multimedia and XML DBMSs, for example. Traditionally, performance evaluation is achieved with the use of benchmarks. While the ability to adapt to changes in access patterns is critical to database performance, none of the existing benchmarks designed for ODBMSs incorporate the possibility of change in the pattern of object access. However, in real life, almost all applications do not always access the same objects in the same order repeatedly. Furthermore, none of the numerous studies regarding dynamic object clustering contain any indication of how these algorithms are likely to perform in a dynamic setting.

In contrast to the TPC benchmarks [13] that aim to provide *standardised* means of comparing systems, we have designed the *Dynamic Object Evaluation Framework* (DOEF) to provide a means to explore the performance of databases under *different* styles of access pattern change.

DOEF contains a set of protocols which in turn define a set of styles of access pattern change. DOEF by no means has exhausted all possible styles of access pattern change. However, it makes the first attempt at exploring the issue of evaluating ODBMSs in general and dynamic clustering algorithms in particular, with respect to changing query profiles.

DOEF is built on top of the *Object Clustering Benchmark* (OCB) [8], which is a generic benchmark that is able to simulate the behavior of the *de facto* standards in object-oriented benchmarking (namely OO1 [5], HyperModel [1], and OO7 [3]). DOEF uses both the database built from the rich schema of OCB and the operations offered by OCB. DOEF is placed into a non-intrusive part of OCB, thus making it clean and easy to implement on top of an existing OCB implementation. Furthermore, we have designed DOEF to be open and fully extensible. First, DOEF's design allows new styles of change to be easily incorporated. Second, OCB's generic model can be implemented within an object-relational system and most of its operations are relevant for such a system. Hence, DOEF can also be used in the object-relational context.

To illustrate the capabilities of DOEF, we benchmarked four state of the art dynamic clustering algorithms. There are three reasons for choosing to test the effectiveness of DOEF using dynamic clustering algorithms: “ever since the early days of object database management systems, clustering has proven to be one of the most effective performance enhancement techniques” [10]; the performance of dynamic clustering algorithms are very sensitive to changing access patterns; and despite this sensitivity, no previous attempt has been made to benchmark these algorithms in this way.

This paper makes two key contributions: (1) it proposes the first evaluation framework that allows ODBMS and associated optimisation techniques to be evaluated in a dynamic environment; (2) it presents the first performance evaluation experiments of dynamic clustering algorithms in a dynamic environment.

The remainder of this paper is organised as follows: Section 2 describes the DOEF framework in detail, we present and discuss experimental results achieved with DOEF in Section 3, and finally conclude the paper and provide future research directions in Section 4.

2 Specification of DOEF

2.1 Dynamic Framework

We start by giving an example scenario that the framework can mimic. Suppose we are modeling an on-line book store in which certain groups of books are popular at certain times. For example, travel guides to Australia during the 2000 Olympics may have been very popular. However, once the Olympics is over, these books may suddenly or gradually become less popular. Once the desired book has been selected, information relating to the book may be required. Example required information includes customer reviews of the book, excerpts from the book, picture of the cover, etc. In an ODBMS, this information is stored

as objects referenced by the selected object (book), thus retrieving the related information is translated into an object graph navigation with the traversal root being the selected object (book). After looking at the related information for the selected book, the user may choose to look at another book by the same author. When information relating to the newly selected book is requested, the newly selected object (book) becomes the root of a new object graph traversal. We now give an overview of the five main steps of the dynamic framework and in the process show how the above example scenario fits in.

1. **H-region parameters specification:** In this step we divide the database into regions of homogeneous access probability (H-regions). In our example, each H-region represents a different group of books, each group having its own probability of access.
2. **Workload specification:** H-regions are responsible for assigning access probability to objects. However, H-regions do not dictate what to do after an object has been selected. We term the selected objects *workload root*, or simply *root*. In this step, we select the type of workload to execute after selecting the root from those defined in OCB. In our example, the selected workload is an object graph traversal from the selected book to information related to the selected book, e.g., an excerpt.
3. **Regional protocol specification:** Regional protocols use H-regions to accomplish access pattern change. Different styles of access pattern change can be accomplished by changing the H-region parameter values with time. For example, a regional protocol may initially define one H-region with high access probability, while the remaining H-regions are assigned low access probabilities. After a certain time interval, a different H-region may become the high access probability region. This, when translated to the book store example, is similar to Australian travel books becoming less popular after the 2000 Olympics ends.
4. **Dependency protocol specification:** Dependency protocols allow us to specify a relationship between the currently selected root and the next root. In our example, this is reflected in the customer deciding to select a book which is by the same author as the previously selected book.
5. **Regional and dependency protocol integration specification:** In this step, regional and dependency protocols are integrated to model changes in dependency between successive roots. An example is a customer using our on-line book store, who selects a book of interest, and then is confronted with a list of *currently* popular books by the same author. The customer then selects one of the listed books (modeled by dependency protocol). The set of *currently* popular books by the same author may change with time (modeled by regional protocol).

2.2 H-regions

H-regions are database regions of homogeneous access probability. The parameters that define H-regions are listed below.

- *HR_SIZE*: Size of the H-region (fraction of the database size).
- *INIT_PROB_W*: Initial probability *weight* assigned to the region. The actual probability is equal to this probability weight divided by the sum of all probability weights.
- *LOWEST_PROB_W*: Lowest probability weight the region can go down to.
- *HIGHEST_PROB_W*: Highest probability weight the region can go up to.
- *PROB_W_INCR_SIZE*: Amount by which the probability weight of the region increases or decreases when change is requested.
- *OBJECT_ASSIGN_METHOD*: Determines the way objects are assigned into the region. *Random* selection picks objects randomly from anywhere in the database. By *class* selection first sorts objects by class ID and then picks the first N objects (in sorted order), where N is the number of objects allocated to the H-region.
- *INIT_DIR*: Initial direction that the probability weight increment moves in.

2.3 Regional Protocols

Regional protocols simulate access pattern change by first initialising the parameters of every H-region, and then periodically changing the parameter values in certain predefined ways. This paper documents three styles of regional change. For every regional protocol, a user defined parameter H is used to control the rate at which access pattern changes. More precisely, H is defined as one divided by the number of transactions executed between each change of access pattern. Three regional protocols are listed below:

- **Moving Window of Change Protocol:** This regional protocol simulates sudden changes in access pattern. In our on-line book store, this is translated to books suddenly becoming popular due to some event (e.g., a TV show). Once the event passes, the books become unpopular very fast. This style of change is accomplished by moving a window through the database. The objects in the window have a much higher probability of being chosen as root when compared to the remainder of the database. This is done by breaking up the database into N H-regions of equal size. Then, one H-region is first initialised to be the hot region (i.e., a region with high probability of reference), and after a certain number of root selections, a different H-region becomes the hot region.
- **Gradual Moving Window of Change Protocol:** This protocol is similar to the previous one, but the hot region cools down gradually instead of suddenly. The cold regions also heat up gradually as the window is moved onto them. This tests the dynamic clustering algorithm’s ability to adapt to a more moderate style of change. In our book store example, this style of change may depict travel guides to Australia gradually becoming less popular after the Sydney 2000 Olympics. As a consequence, travel guides to other countries may gradually become more popular. Gradual changes of heat may be more common in the real world. This protocol is implemented in the same way as the previous protocol except the H-region that the window (called

the hot region in the previous protocol) moves *into* gradually heats up and the H-region that the window moves *from* gradually cools down.

- **Cycles of Change Protocol:** This style of change mimics something like a bank where customers in the morning tend to be of one type and in the afternoon of another type. This, when repeated, creates a cycle of change. This is done by break up the database into three H-regions. The first two H-regions represent objects going through the cycle of change. The third H-region represent the remaining unchanged part of the database. The first two H-regions alternates at being the hot region.

2.4 Dependency Protocols

There are many scenarios in which a person executes a query and then decides to execute another query based on the results of the first query, thus establishing a dependency between the two queries. In this paper, we have specified four dependency protocols. All four protocols functions by finding a set of candidate objects that maybe used as the next root. Then a random function is used to select one object out of the candidate set. The selected object is the next root. An example random function is a skewed random function that selects a certain subset of candidate objects with a higher probability than others. The four dependency protocols are listed below:

- **Random Selection Protocol:** This method simply uses some random function to select the current root. This protocol mimics a person starting a completely new query after finishing the previous one.
- **By Reference Selection Protocol:** The current root is chosen to be an object referenced by the previous root. An example of this protocol in our on-line book store scenario is a person having finished with a selected book, who then decides to look at the next book in the series.
- **Traversed Objects Selection Protocol:** The current root is selected from the set of objects that were referenced in the previous traversal. An example is a customer requesting in a first query a list of books along with their author and publisher, who then decides to read an excerpt from one of the books listed.
- **Same Class Selection Protocol:** The currently selected root must belong to the same class as the previous root. Root selection is further restricted to a subset of objects of the class. The subset is chosen by a function that takes the previous root as a parameter. That is, the subset chosen dependent on the previous root object. An example of this protocol is a customer deciding to select a book from our on-line book store which is by the same author as the previous selected book. In this case, the same class selection function returns books by the same author.

Hybrid Setting. The hybrid setting allows an experiment to use a mixture of the dependency protocols outlined above. Its use is important since it simulates a user starting a fresh random query after having followed a few dependencies.

Thus, the hybrid setting is implemented in two phases. The first *randomisation phase* uses the random selection protocol to randomly select a root. In the second *dependency phase*, one of the dependency protocols outlined in the previous section is used to select the next root. R iterations of the second phase are repeated before going back to the first phase. The two phases are repeated continuously.

2.5 Integration of Regional and Dependency Protocols

Dependency protocols model user behavior. Since user behavior can change with time, dependency protocols should also be able to change with time. The integration of regional and dependency protocols allows us to simulate changes in the dependency between successive root selections. This is easily accomplished by exploiting the dependency protocols' property of returning a candidate set of objects when given a particular previous root. Up to now, the next root is selected from the candidate set by the use of a random function. Instead of using the random function, we partition the candidate set using H-regions and then apply regional protocols on these H-regions.

3 Experimental Results

3.1 Experimental Setup

We used DOEF to compare the performance of four state of the art dynamic clustering algorithms: Dynamic, Statistical, and Tunable Clustering (DSTC) [2], Detection & Reclustering of Objects (DRO) [6], dynamic Probability Ranking Principle (PRP) [12], and dynamic Graph Partitioning (GP) [12]. The aim of dynamic clustering is to automatically place objects that are likely to be accessed together in the near future in the same disk page, thereby reducing the number of I/O. The four clustering techniques have been parameterized for the same behaviour and best performance.

We chose simulation for these experiments, principally because it allows rapid development and testing of a large number of dynamic clustering algorithms (all previous dynamic clustering papers compared at most two algorithms). The experiments were conducted on the Virtual Object-Oriented Database discrete-event simulator (VOODB) [7]. Its purpose is to allow performance evaluations of OODBs in general, and optimisation methods like clustering in particular. VOODB has been validated for real-world OODBs in a variety of situations. The VOODB parameter values we used are depicted in Table 1 (a).

Since DOEF uses the OCB database and operations, it is important for us to document the OCB settings used for these experiments (Table 1 (b)). The size of the objects used varied from 50 to 1600 bytes, with an average of 233 bytes. A total of 100,000 objects were generated for a total database size of 23.3 MB. Although this is a small database size, we also used a small buffer size (4 MB) to keep the database to buffer size ratio large. Clustering algorithm performance is indeed more sensitive to database to buffer size ratio than database size alone.

The operation used for all the experiments was a simple, depth-first traversal of depth 2. We chose this simple traversal because it is the only one that always accesses the same set of objects given a particular root. This establishes a direct relationship between varying root selection and changes in access pattern. Each experiment involved executing 10,000 transactions. The main DOEF parameter settings used in this study are shown in Table 2. These DOEF settings are common to all experiments in this paper. The *HR_SIZE* setting of 0.003 creates a hot region about 3% the size of the database. This fact was verified from statistical analysis of the trace generated. The *HIGHEST_PROB_W* setting of 0.8 and *LOWEST_PROB_W* setting of 0.0006 produce a hot region with 80% probability of reference, the remaining cold regions having a combined reference probability of 20%. These settings are chosen to represent typical database application behaviour [11, 4, 9].

(a) VOODB parameters		(b) OCB parameters	
Parameter Description	Value	Parameter Description	Value
System class	Centralized	Number of classes	50
Disk page size	4096 bytes	Maximum number of references, per class	10
Buffer size	4 MB	Instances base size, per class	50
Buffer replacement policy	LRU-1	Total number of objects	100000
Pre-fetching policy	None	Number of reference types	4
Multiprogramming level	1	Reference types distribution	Uniform
Number of users	1	Class reference distribution	Uniform
Object initial placement	Sequential	Objects in classes distribution	Uniform
		Objects references distribution	Uniform

Table 1. VOODB and OCB parameters

Parameter Name	Value
<i>HR_SIZE</i>	0.003
<i>HIGHEST_PROB_W</i>	0.80
<i>LOWEST_PROB_W</i>	0.0006
<i>PROB_W_INCR_SIZE</i>	0.02
<i>OBJECT_ASSIGN_METHOD</i>	Random object assignment

Table 2. DOEF parameters

As we discuss the results of these experiments, we focus our discussion on the relative ability of each algorithm to adapt to changes in access pattern, i.e., as rate of access pattern change increases, we seek to know which algorithm exhibits more rapid performance deterioration. This contrasts from discussing which algorithm gives the best absolute performance. All the results presented here are in terms of total I/O (transaction I/O plus clustering I/O).

3.2 Moving and Gradual Moving Regional Experiments

In these experiments, we tested the dynamic clustering algorithms' ability to adapt to changes in access pattern by varying the rate of access pattern change

(parameter H). The results of these experiments (Figure 1) induce three main conclusions. First, when rate of access pattern change is small (H lower than 0.0006), all algorithms show similar performance trends. Second, when the more vigorous style of change is applied (Figure 1 (a)), all dynamic clustering algorithms' performance quickly degrades to worse than no clustering. Third, when access pattern change is very vigorous (H greater than 0.0006), DRO, GP, and PRP show a better performance trend, implying these algorithms are more robust to access pattern change. This is because these algorithms choose only a relatively few pages (the worst clustered) to re-cluster. This leads to greater robustness. We term this flexible conservative re-clustering. In contrast, DSTC re-clusters a page even when there is only small potential gain. This explains DSTC's poor performance when compared to the other algorithms.

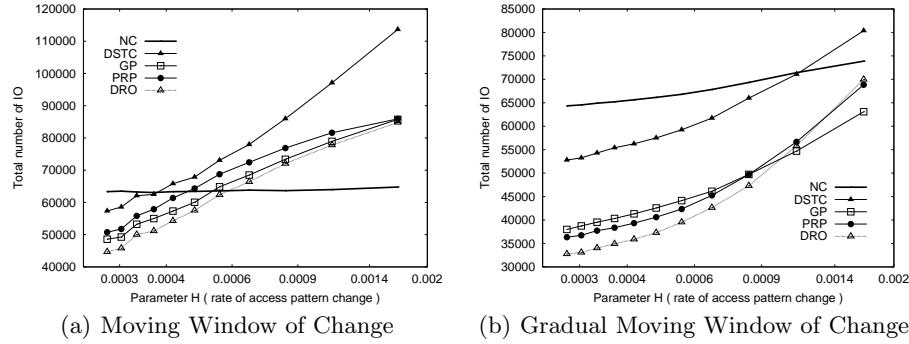


Fig. 1. Regional dependency results

3.3 Moving and Gradual Moving By Reference Experiments

These experiments, used the integrated regional dependency protocol method outlined in Section 2.5 to integrate *by reference dependency* with the *moving* and *gradual moving window of change regional* protocols. We also used the hybrid dependency setting detailed in Section 2.4. The random function we used in the first phase partitioned the database into one hot (3% database size and 80% probability of reference, which represents typical database application behaviour) and one cold region. The results for these experiments are shown on Figure 2. In the moving window of change results (Figure 2 (a)), DRO, GP and, PRP were again more robust to changes in access pattern than DSTC. However, in contrast to the previous experiment, DRO, GP, and PRP never perform worse than NC by much, even when parameter H is 1 (access pattern changes after every transaction). The reason is the cooling and heating of references is a milder form of access pattern change than the pure moving window of change regional protocol of the previous experiment. As in the previous experiment, all dynamic clustering algorithm show approximately the same performance trend for the gradual moving window of change results (Figure 2 (b)).

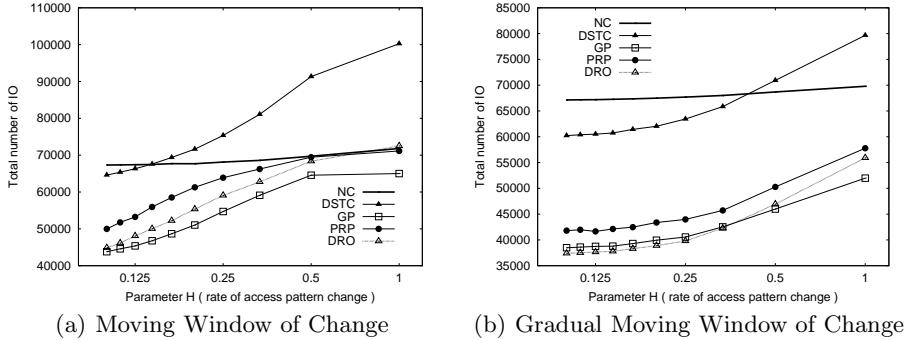


Fig. 2. S-reference dependency results

4 Conclusion

In this paper we presented a new framework for object benchmarking, DOEF, which allows ODBMSs' designers and users to test the performances of a given system in a dynamic setting. This is an important contribution since almost all real world applications exhibit access pattern changes, but no existing benchmark attempt to model this behavior.

We have designed DOEF to be extensible along two axes. First, new styles of access pattern change can be defined, through the definition of H-regions. We encourage other researchers or users to extend DOEF by making the DOEF code freely available for download¹. Second, we can apply the concepts developed in this paper to object-relational databases. This is made easier by the fact OCB (the layer below DOEF) can be easily adapted to the object-relational context².

Experimental results have demonstrated DOEF's ability to meet our objective of exploring the performance of databases within the context of changing patterns of data access. Two new insight were gained: dynamic clustering algorithms can cope with moderate levels of access pattern change but performance rapidly degrades to be worse than no clustering when vigorous styles of access pattern change is applied; and flexible conservative re-clustering is the key in determining a clustering algorithm's ability to adapt to changes in access pattern.

This study opens several research perspectives. The first one concerns the exploitation of DOEF to keep on aquiring knowledge about the dynamic behavior of various ODBMSs. Second, adapting OCB and DOEF to the object-relational model will enable performance comparison of object-relational DBMSs. Since OCB's schema can be directly implemented within an object-relational system, this would only involve adapting existing and proposing new OCB operations relevant for such a system. Lastly, the effectiveness of DOEF at evaluating other

¹ <http://eric.univ-lyon2.fr/~jdarmont/download/docb-voodb.tar.gz>

² Even if extensions would be required, such as abstract data types or nested tables.

aspects of database performance could be explored. Optimisation techniques, such as buffering and prefetching could also be evaluated.

References

1. ANDERSON, T., BERRE, A., MALLISON, M., PORTER, H., AND SCHEIDER, B. The HyperModel benchmark. In *2nd International Conference on Extending Database Technology (EDBT)* (March 1990), vol. 416 of *LNCS*, pp. 317–331.
2. BULLAT, F., AND SCHNEIDER, M. Dynamic clustering in object databases exploiting effective use of relationships between objects. In *10th European Conference on Object Oriented Programming (ECOOP)* (1996), vol. 1222 of *LNCS*, pp. 344–365.
3. CAREY, M., DEWITT, D., AND NAUGHTON, J. The OO7 benchmark. *ACM SIGMOD International Conference on Management of Data* (May 1993), 12–21.
4. CAREY, M. J., FRANKLIN, M. J., LIVNY, M., AND SHEKITA, E. J. Data caching tradeoffs in client-server DBMS architectures. In *ACM SIGMOD International Conference on Management of Data* (1991), pp. 357–366.
5. CATTELL, R. An engineering database benchmark. In *The Benchmark Handbook for Database Transaction Processing Systems* (1991), pp. 247–281.
6. DARMONT, J., FROMANTIN, C., REGNIER, S., GRUENWALD, L., AND SCHNEIDER, M. Dynamic clustering in object-oriented databases: An advocacy for simplicity. In *International Symposium on Object and Databases* (June 2000), vol. 1944 of *LNCS*, pp. 71–85.
7. DARMONT, J., AND SCHNEIDER, M. VOODB: A generic discrete-event random simulation model to evaluate the performances of OODBs. In *25th International Conference on Very Large Data Bases (VLDB)* (September 1999), pp. 254–265.
8. DARMONT, J., AND SCHNEIDER, M. Benchmarking OODBs with a generic tool. *Journal of Database Management* 11, 3 (Jul-Sept 2000), 16–27.
9. FRANKLIN, M. J., CAREY, M. J., AND LIVNY, M. Local disk caching for client-server database systems. In *19th International Conference on Very Large Data Bases (VLDB)* (1993), R. Agrawal, S. Baker, and D. A. Bell, Eds., pp. 641–655.
10. GERLHOF, C., KEMPER, A., AND MOERKOTTE, G. On the cost of monitoring and reorganization of object bases for clustering. *ACM SIGMOD Record* 25 (1996), 28–33.
11. GRAY, J., AND PUTZOLU, G. R. The 5 minute rule for trading memory for disk accesses and the 10 byte rule for trading memory for cpu time. In *ACM SIGMOD International Conference on Management of Data* (1987), pp. 395–398.
12. HE, Z., MARQUEZ, A., AND BLACKBURN, S. Opportunistic prioritised clustering framework (OPCF). In *International Symposium on Object and Databases* (June 2000), vol. 1944 of *LNCS*, pp. 86–100.
13. TRANSACTION PROCESSING PERFORMANCE COUNCIL. TPC home page. <http://www.tpc.org>, 2002.

